

CONFIGURATION OF FIBRE CHANNEL SAN PATH

TECHNICAL FIELD

- [1] The present invention relates generally to fibre channel storage area networks and, in particular, to configuring a path between a port in a source node and a port in a target node.

BACKGROUND ART

- [2] A storage area network ("SAN") permits the distribution of components, such as servers, storage devices and storage controllers, to provide a customer with access to a large storage pool. Fig. 1 is a block diagram of a generic SAN 100. Controllers 102 may be shared among multiple servers 104. In one type of SAN, the components are interconnected using the Fibre Channel protocol. The components (also referred to as nodes) may be interconnected directly (a point-to-point topology), in a ring (an arbitrated loop topology) or through switches (a switched fabric topology). The term "switch" is used herein generically to refer to switches, directors, hubs, bridges and the like. Using switches, data may be delivered across the fabric 106 from any source (or initiator) node to any destination (or target) node.
- [3] Each node is given a unique 8-byte address (a "world wide node name" or WWNN) which is assigned by the manufacturer in accordance with an established IEEE standard. Each port in a node is also given a unique 8-byte address (a "world wide port name" or WWPN) by the manufacturer.
- [4] In part to reduce the space in a channel frame header required to address a particular port and thus speed up routing of data, another port addressing scheme was developed. After a connection has been established between a node and a switch, each port in the node is assigned a unique (within the fabric) address (destination_id) by a controlling switch using a name server. The switch maintains a translation between the port address and the WWN address and can address any port in the fabric using the destination_id. Additionally, each slot of a node has a "name" (interface_id) by which the slot may be physically identified.

- [5] However, because there may be no correlation between a port address and the physical slot location of the port, a human administrator may find it difficult to determine to which port of a node a cable should be connected. Consequently, a need exists for easily determining which physical slot on a node corresponds with a port address.

SUMMARY OF THE INVENTION

- [6] The present invention provides method, system and computer program product to configure a path between nodes through a fabric in a fibre channel storage area network (SAN). A node name is provided for a target node on the SAN and a port name is provided for each port in the target node. A relationship is established in a data structure between each port name and a slot in which each port is physically located in the target node. Prior to configuring a path between a source node and a port in the target node, the ports which are physically connected and logged in to the fabric are identified by port address. The port names corresponding to the port addresses are then identified. The port names are used to generate interface_ids of the ports corresponding to the physical slots in which the ports are located. A data structure is created to maintain the relationship between interface_ids and port names. To establish a path connection between a port in the source node and a selected port in the target node, the node name of the target node and the interface_id of the selected port are input to the data structure and the port name of the selected port is output. The port name is then used to obtain the address of the selected port and an I/O session path opened between the source and target nodes.
- [7] The name of the target node may comprise the world wide node name of the node and name of the selected port may comprise the world wide port name of the port. The address of the selected port may be the destination_id of the port, assigned by a name server on the fabric.

BRIEF DESCRIPTION OF THE DRAWINGS

- [8] Fig. 1 is a block diagram of a generic storage area network;

- [9] Fig. 2 is a block diagram of a portion of a storage area network in which the present invention may be implemented; and
- [10] Fig. 3 is a flow chart of an embodiment of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

- [11] Fig. 2 is a block diagram of a portion 200 of a storage area network in which the present invention may be implemented. A source or initiator node 210 and a destination or target node 220 are coupled through a switch 230. For example, in a peer-to-peer remote copy (PRPC) environment, the source node 210 may be a primary storage controller and the target node 220 may be a secondary storage controller. Each of the components 210, 220 and 230 has a manufacturer-assigned unique world wide node name (WWNN). The source node 210 includes two ports 212 and 214; the target node includes two ports 222 and 224; and the switch 230 includes four ports 232, 234, 236 and 238. Each of the ports 212, 214, 222, 224, 232, 234, 236 and 238 has a manufacturer-assigned unique world wide port name (WWPN). As will be appreciated, nodes may include more ports than are illustrated. Channel links 240, 242, 244 and 246 interconnect the ports. Additionally, each port 212, 214, 222 and 224 in the two nodes 210 and 220 is physically located in a slot in a node and has a switch-assigned destination_id.
- [12] A user attempting to physically connect two of the components, for example the source node 210 with the switch 230, needs to know the slot number of each of the ports on the components. Referring to the flow chart of Fig. 3, in accordance with the present invention, one of the bytes of the WWNN of one of the components (for example, the source node 210) is modified to indicate the slot number in which the connected port is located to provide the WWPN (step 302). In certain storage controllers manufactured by IBM®, byte 5 of the WWNN is set to a value of C0. The slot number will replace this value. For example, if the port is located in slot number 3, byte 5 of the WWNN will be changed from C0 to C3 to thus generate the WWPN.
- [13] Next, a data structure, such as a table or algorithm, is created and stored. The data structure links or cross-references the interface_id of the slot in which a port is located with the WWPN of the port (step 304). In this manner, a relationship is

established between the WWPN and the physical slot (designated by the interface_id).

- [14] In order to establish a link between two nodes, it is necessary for the source node to first identify which ports of the target node 220 are logged into the fabric; that is, to determine what physical connections exist between the two nodes. The source node 210 uses a name server on the switch 230 to transmit a query onto the fabric based on the WWNNs of the source and target nodes (step 306). The source node 210 receives back a list of port addresses (destination_ids) which are physically connected to and logged in to the fabric under the specified target WWNN (step 308). In response to a further query (step 310), the name server next returns the WWPN of the connected ports (step 312). Referencing the data structure, the interface_id of each connected port (on the target node) is determined (step 314).
- [15] When it is desired to establish a path from the source node 210 to the target node 220, the source node 210 again accesses the data structure to cross reference the interface_id and WWNN of the selected target port to determine the WWPN (step 316). Using the WWPN, the source node 210 queries the fabric to obtain the actual port address (destination_id) of the target port (step 318). The returned port address (step 320) is then used to perform a logon and a session is opened for I/O activities between the source 210 and target 220 nodes (step 322).
- [16] In the event that the target hardware is replaced or the configuration of the SAN changes, the port name (WWPN) and/or address (destination_id) will also change. However, when the present invention is implemented, a path remains established independent of the physical changes.
- [17] The objects of the invention have been fully realized through the embodiments disclosed herein. Those skilled in the art will appreciate that the various aspects of the invention may be achieved through different embodiments without departing from the essential function of the invention. The particular embodiments are illustrative and not meant to limit the scope of the invention as set forth in the following claims.